

IN THE ABSTRACT

A system and method of allocating a resource to a service request. Servers are virtually assigned through a hierarchy including groups of servers and super groups of the groups of servers. Each service request type is associated with a service index that determines both a super group and a load balancing policy for that service request type. A load balancing policy may be applied to select a group of servers among several server groups associated with the super group. Another load balancing policy is applied to the selected group of servers to select a server to handle the service request. The hierarchical data structure allows servers to be configured into overlappable, arbitrary subsets that can address service requests matching configured content rules. Load balancing policies can be selected by service request type.